Bee vision

A dimensionality reduction approach

Sashank Pisupati

Mentor: Prof. Amitabh Mukerjee

Abstract

Bees rely heavily on visual cues to locate themselves, and combine these with scent cues to learn about nectar sources through reinforcement learning[3][7]. From neural data, we know that the neurons that perform this learning must receive input that is low dimensional, in order to couple it reliably to motor feedback[2]. This low dimensional mapping of the high dimensional visual space may allow the bee to interpolate views, plan its path and learn about nectar sources.[8][4][6]

From the very different context of robot path planning/visuomotor discovery, we know that a robot can use images of its own parts in various configurations to infer its degrees of freedom, since the manifold on which the images lie is of the same dimension as its degrees of freedom[1]

This project then asks the question: Can a bee foraging in an environment, discover the low dimensional visual subspace that is meaningful for its perception through nonlinear dimensionality reduction, and will the dimension of this space be representative of its degrees of freedom of motion? If so, then the bee can use this low dimension to infer its location and configuration in space, and use this as input to its odor/nectar sensitive neurons, and learn salient locations/plan foraging paths.

Introduction

Bees are one of the most well studied insect models of perception, since they rely so heavily on visual information for navigation and motor feedback. Bees can use visual cues to locate themselves in the environment, to learn to distinguish good sources of nectar from bad ones, and most crucially, to control flight. Bees can use just optic flow velocity of their visual scene to modulate flight speed, perform grazing landings, negotiate tiny slits etc.[6]

Neural studies have unearthed single neurons receiving input from the visual areas of the bee, that are sensitive to **low-dimensional** variations in visual input, such as direction of optic motion, that are strongly coupled to motor output. These neurons are capable of providing motor feedback, performing reinforcement learning to plan motor paths on the basis of nectar availability[3] and so on. For example, Srinivasan et. al, while describing the bee's optomotor response, mention how movement sensitive neurons in the bee simply need to reliably detect the direction of motion in order to correct their flight path[8].

In 1995, Montague, Sejnowski et. al demonstrated a computational model of such a reinforcement learning neuron, that took percentages of different colored flowers in its environment as low-dimensional input, and used a nectar reward to bias its actions while foraging.(Shown below from Montague et. al, the virtual bee in its environment, and the neuron P that biases the path of the foraging bee). Such neurons have been observed to play a role in risk aversion, and other reward related actions.

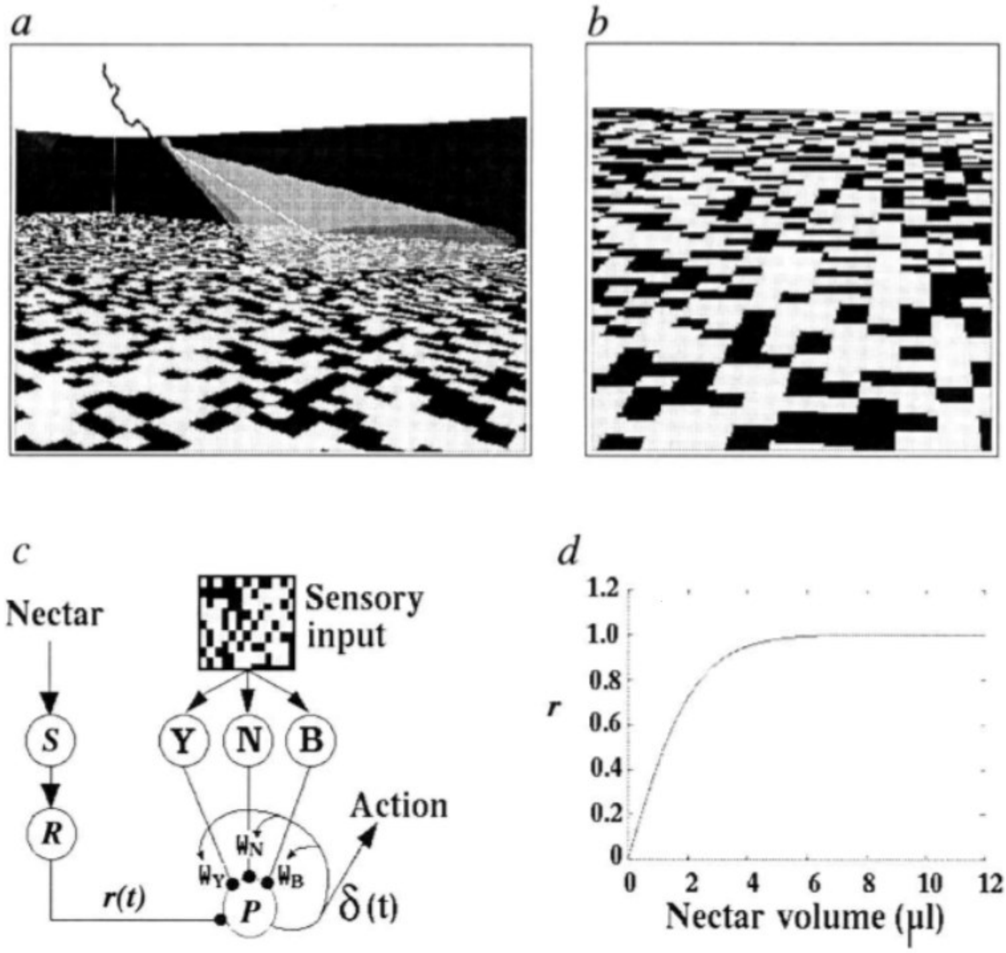

Dimensionality reduction: Theoretic considerations

Why dimensionality reduction? The visual input to a bee's eye is in itself very high dimensional. But as mentioned before, perceptually meaningful information is much lower dimensional and can be used by the aforementioned neurons.

In 2005, Bialek, Steveninck and Ruyter looked at fly visual neurons, and hypothesised that feature selective neurons are *only sensitive* to low-dimensional sub-spaces of input from the visual stream, describing a non-linear mapping from the high dimensional visual space to this sub-space. Moreover, recent studies have shown that bees are capable of complex visual tasks like **image interpolation** between different views of an object [9]

Given this knowledge, it is reasonable to hypothesise that a non-linear dimensionality reduction of the visual input could be taking place in the visual system, enabling such mappings/computations.

How is it useful?
Now, if the visual system receives certain kinds of images owing to the motion of the agent through the environment, then the low-dimensional manifold on which the images lie must convey information about this motion. In the context of robot learning, it has been shown that if a robot tries to discover the low dimensional manifold of images of its own arms, then the dimension of this manifold will correspond to the number of degrees of freedom of its arms (Mukherjee et. al, [1]). Hence, using this low dimensional manifold, the robot can infer the position of its arms and use this for motion planning i.e. it is useful for visuo-motor feedback.

The hypothesis of this project is that, similarly, if a virtual bee moves around an environment, receiving **different views of the environment** owing to its motion, then, on performing a non-linear dimensionality reduction on the images, the **dimension** of the low dimensional manifold on which the images lie will correspond to the degrees of freedom of its motion.

This would then let the bee infer, on the basis of visual stimulus alone, its configuration or location while foraging.

Methods

The eye and the environment of the virtual bee was simulated using Andy Giger's java simulation, "B-EYE". This simulates the image projected onto the photoreceptors of the bee's compound eye, as the "bee" (represented by the cursor) moves around in three-dimensional space around a flat image. The XY movement is performed by moving around the cursor while the Z-movement is performed by zooming in or out with right/left clicks.

The images below are screenshots of the B-EYE simulation, with example trajectories of the virtual bee as it was manually moved around the environment. The first image shows a 2-D trajectory, while the second image shows a 3-D trajectory with different colors representing different levels of zoom (Z-movement). As this was being done, a 300x300 square portion of the bee visual array was continuously captured using the "screenr" screen capture software, and images were extracted from this at a rate of 10fps (sample images are displayed below, alongside the experimental setup)

Then, a nonlinear dimensionality reduction was performed on the images using Isomap, an implementation of the Dijkstra algorithm, to discover the low dimensional manifold on which the images lie. The dimension of this manifold as well as the low-dimensional embedding of the images on it was then plotted.

Results

The following results were obtained for various trajectories in 2d and 3d. Many of the plots did seem to be representative of the degrees of freedom of motion of the trajectory, although some, especially ones with very poor sampling of the environment, showed higher dimensions than their DoFs (typically one dimension higher). Also, when tested on environments with repetitive patterns, and rotated versions of the same (eg. Image 2 in the set of testable images) some interesting results emerged.

2d trajectories: the favourable

The following graphs show trials involving 2d motion in XY. The scree plot in all these cases shows a sharp kink at dimension=2 (change from a steep slope to a much lesser slope) indicating that two dimensions explain the visual data. In other words, the bee might infer that it is moving about in two dimensions. The low dimensional embeddings show the distribution of images on the 2d manifold, based on the positions of which the bee can get a two-dimensional description of its position (i.e. its XY configuration).

The first two images are of a flower, and the third image is of a striped pattern, where the stripes where along the vertical, but being small circular patterns, their distortions on the eye varied along the horizontal (possibly conveying two degrees of motion rather than just movement of stripes along the vertical) and the fourth image is of a wall painting around which the bee is moving in 2d.

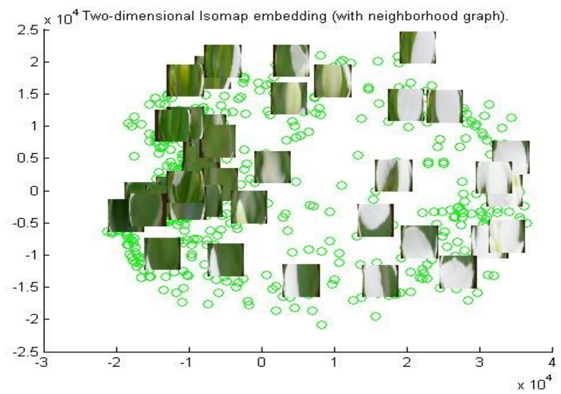

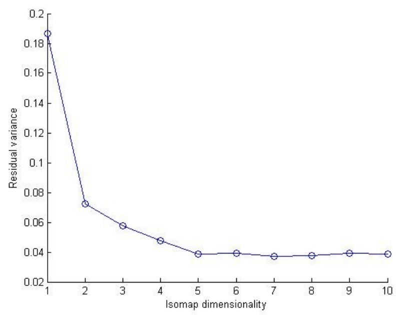

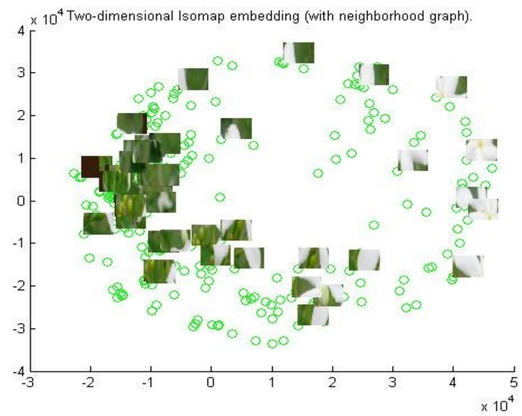

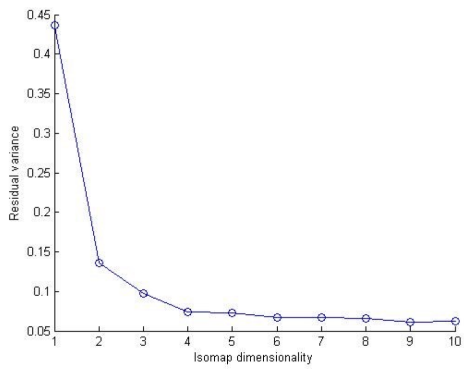

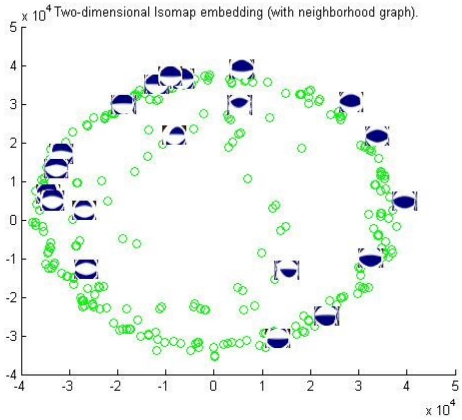

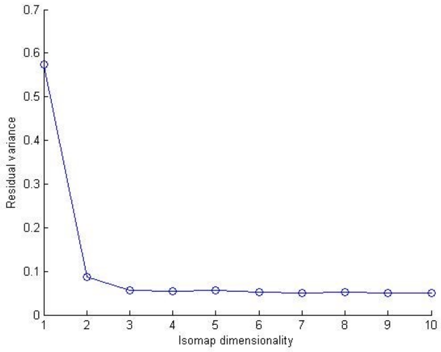

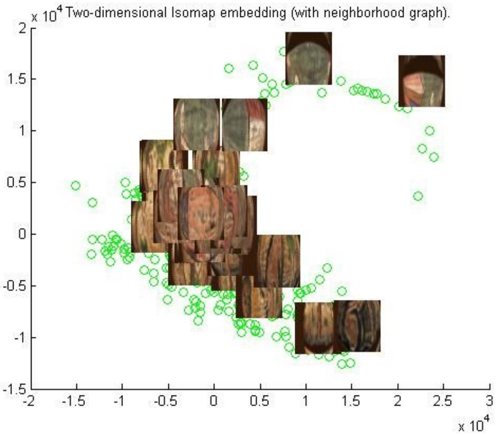

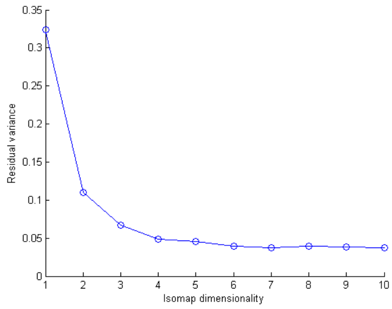

3d trajectories: the favourable

The following graphs depict the dimension plots as well as the embedding for 3d trajectories, that involved zooming in the Z direction. The first two are trajectories with X,Y and Z movement. The third is interesting, because it involved 2-D motion with the second figure in the java program (the set of blue striped circles).

This trajectory was performed very close to the gratings, in such a way that the bee at any point of time gets only stripes of a single grating, but this was done for all six gratings. A possible explanation for the 3-dimensionality is that since the gratings are rotated versions of each other, this might induce the bee to think that it has one extra, rotational degree of freedom.

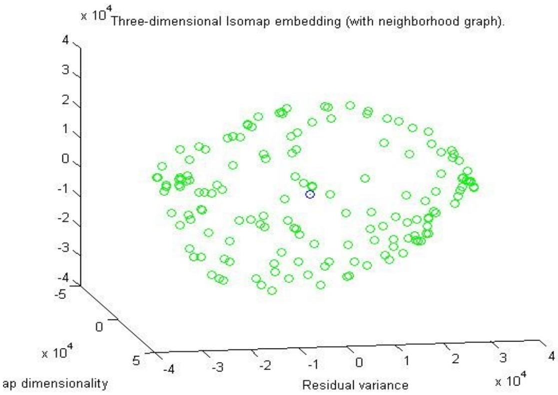

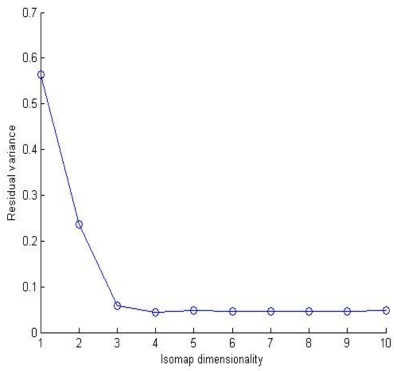

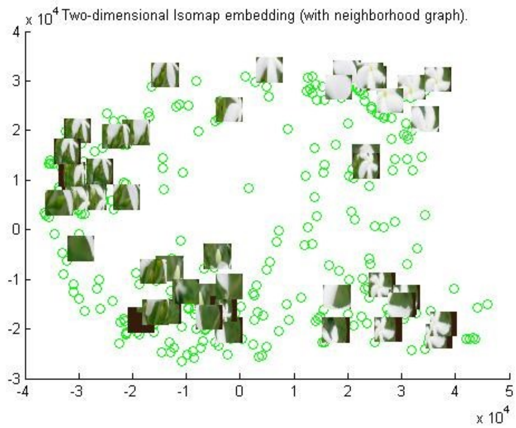

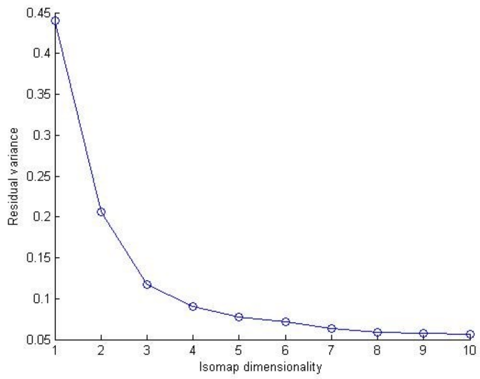

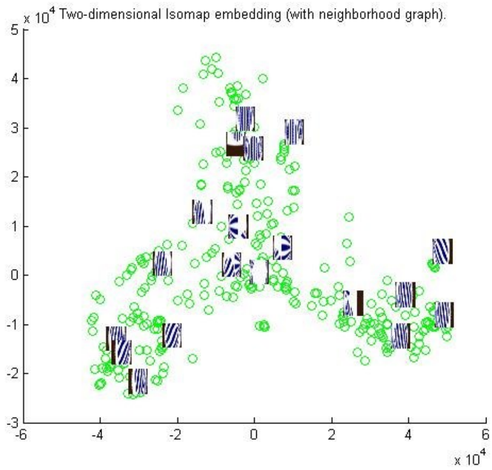

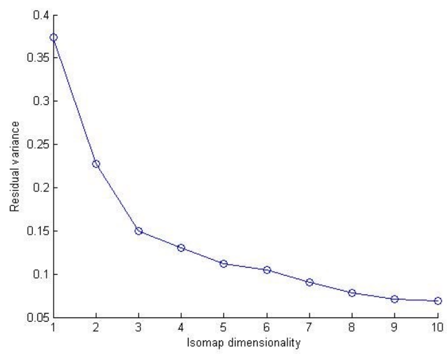

Oddballs: confused bees

The previous data showed quite strong relations between the dimensions and DoFs of movement in the trajectories. The following data, however, did not provide any meaningful inferences. The first row is 2d trajectories and second row is 3d.

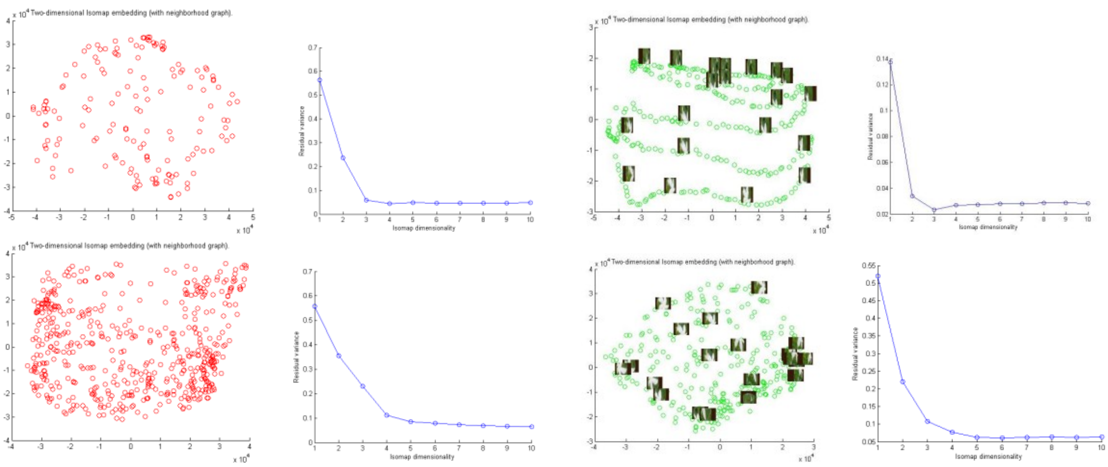

The following plots, however were very interesting, because they once again involved stimulus patterns with poor information along one of their dimensions, i.e. stripes.Such stripe experiments have been often conducted with live bees in an attempt to disambiguate the bees' optomotor coupling. Striped patterns often tend to confuse bees because it may give them the impression of using less degrees of freedom than they are.

The first was an attempt to achieve 2d motion by simply moving along the perpendicular direction of the stripes(1 dimension) but for a variety of stripe orientations(rotation: 2nd dimension).

The second was an attempt in which only two perpendicular stripe patterns were chosen, and 2d motion was performed around them.

The third was a similar attempt for only horizontal and vertical stripes.

All these attempts yielded strange, but interesting manifolds, as well as dimension plots.

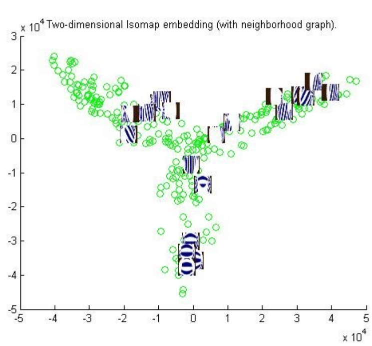

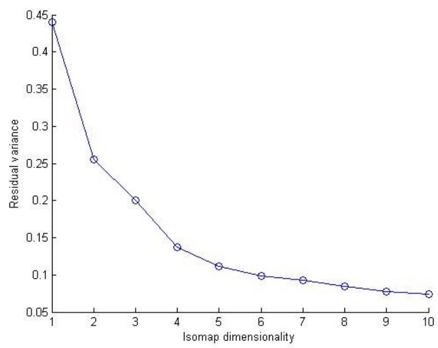

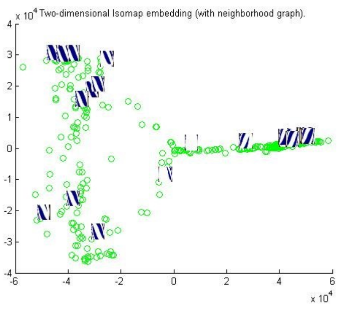

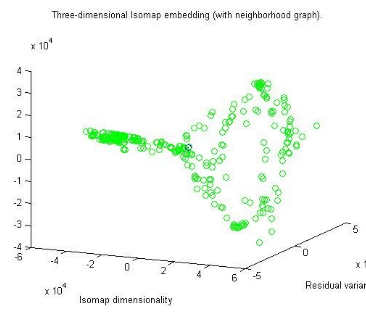

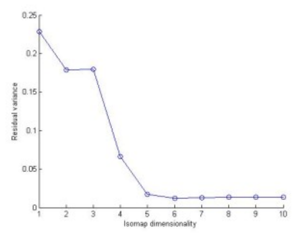

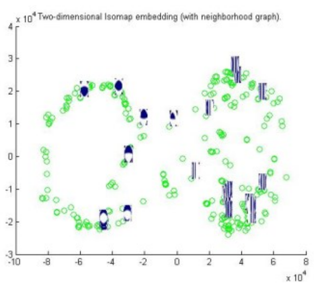

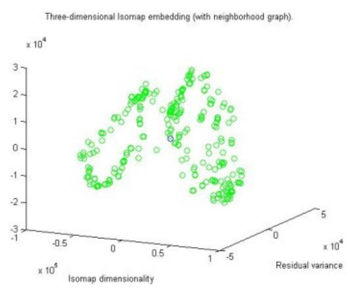

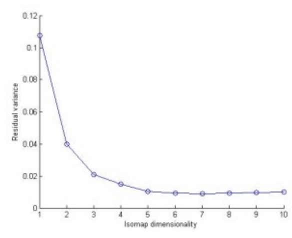

Conclusions

From the majority of results, we could say that the low-dimensional variance in the visual input really is representative of the degrees of freedom of the bee, i.e the ego-motion of the bee through an environment can be reflected in the visual stimulus it gets, quite reliably.

The bee can hence use this low dimensional description to infer its location, and use this information to learn reward-based behaviours.

The next step: Closed loop bees

So far, we have used an open loop bee that is manually made to forage, and gets some passive visual input. The next logical step would be to close the control loop and couple the low dimensional visual input with motor output. Since the dimension matches the degrees of freedom of motion, the bee could hence use the coordinates along these dimensions as inputs to motion planning or reinforcement learning neurons, which could use a neuromodulatory system to build a salience landscape of the bee's environment. Such a bee may be able to actively adjust its scanning of the scene in order to improve the information they get about nectar sources (eg: more dense scanning in salient areas?)

Coming back to the Montague model of foraging, a similar system could be implemented in this virtual bee in a number of ways:

- A reinforcement learning neuron could use these inputs to associate nectar with certain coordinates, thus telling the bee of its location with respect to the closest flower. The output of this could be sent to the motor system, leading to efficient foraging.
- Given the coordinates of two nectar sources, it could try planning a path along one dimension of the manifold, or try to find the shortest path on the manifold between the two.
- A fourth dimension of an odor gradient could be established, and the bee could use its location awareness to try planning a steepest gradient descent path to its destination.

Acknowledgements

I would like to thank Prof. Andy Giger profusely for his wonderful, biologically driven java simulation of the bee's visual array, B-EYE, that was the core experimental setup of this project.

I would like to thank Prof. Amitabh Mukerjee for mentoring me on this project, and for the invaluable support he offered. I would also like to thank M.S. Ram for all his help in implementing the Isomap code and interpreting and fine-tuning the results.

References

Note: Starred citations are the primary references for this report.

Dimensionality reduction:

Low dimensionality and visuo-motor mapping in robots:

[1]* A. Awasthi, S. Sharma, A. Mukerjee: “Learning of motor maps from perception: A dimensionality reduction approach”, available at http://www.cse.iitk.ac.in/users/amit/pub/awasthi-mukerjee-12cogsci-learning-motor- maps-cready.pdf

Low dimensions and feature extraction in insects:

[2]* Bialek, William, and Rob R. van Steveninck. "Features and dimensions: Motion estimation in fly vision." arXiv preprint q-bio/0505003 (2005).

Single neuron computations using low dimensional inputs

[3]* Montague, P. Read, et al. "Bee foraging in uncertain environments using predictive hebbian learning." Nature 377.6551 (1995): 725-728.

[4] Soltoggio, Andrea, et al. "Evolving neuromodulatory topologies for reinforcement learning-like problems." Evolutionary Computation, 2007. CEC 2007. IEEE Congress on. IEEE, 2007.

[5] Niv, Yael, et al. "Evolution of reinforcement learning in foraging bees: A simple explanation for risk averse behavior." Neurocomputing 44 (2002): 951-956.

Bee vision

[6]* Dyer, Adrian G., and Quoc C. Vuong. "Insect brains use image interpolation mechanisms to recognise rotated objects." *PloS one* 3.12 (2008): e4086.

[7] Srinivasan, Mandyam V., and Shaowu Zhang. "Visual motor computations in insects." *Annu. Rev. Neurosci.* 27 (2004): 679-696.

[8] Srinivasan, Mandyam V., Michael Poteser, and Karl Kral. "Motion detection in insect orientation and navigation." *Vision research* 39.16 (1999): 2749-2766.

[9] Paulk, Angelique C., et al. "Visual processing in the central bee brain." *The Journal of Neuroscience* 29.32 (2009): 9987-9999.